

CLAIMS

1. A method for managing a data storage system that includes primary and secondary storage subsystems, including respective first and second non-volatile storage media, the
5 method comprising:

maintaining a record on the secondary storage subsystem, which is predictive of locations to which data are to be written on the primary storage subsystem by a host processor;

10 receiving at the primary storage subsystem, from the host processor, the data to be written to a specified location on the first non-volatile storage media;

if the specified location is not included in the record, sending a message from the primary storage subsystem
15 to the secondary storage subsystem so as to cause the secondary storage subsystem to update the record;

signaling the host processor that the data have been stored in the data storage system responsively to receiving the data and, if the specified location was not included in
20 the record, responsively to receiving an acknowledgment at the primary storage subsystem from the secondary storage subsystem indicating that the record has been updated; and

storing the data in the specified location on both the first and second non-volatile storage media.

25

2. The method according to claim 1, wherein sending the message comprises copying the data synchronously from the primary storage subsystem to the secondary storage subsystem.

3. The method according to claim 2, wherein storing the data comprises, if the specified location is included in the record, copying the data from the primary storage subsystem
5 to the secondary storage subsystem asynchronously, without updating the record with respect to the specified location.

4. The method according to claim 3, wherein copying the data comprises transmitting the data between mutually-remote
10 sites over a communication link between the sites.

5. The method according to claim 3, wherein signaling the host processor comprises, if the specified location is included in the record, indicating to the host processor
15 that the data have been stored without waiting to receive the acknowledgment from the secondary storage subsystem.

6. The method according to claim 1, wherein copying the data comprises creating a mirror on the secondary storage
20 subsystem of the data received by the primary storage subsystem.

7. The method according to claim 6, and comprising, upon occurrence of a failure in the primary storage subsystem,
25 configuring the secondary storage subsystem to serve as the primary storage subsystem so as to receive further data from the host processor to be stored by the data storage system.

8. The method according to claim 6, and comprising, upon recovery of the system from a failure of the primary storage subsystem, conveying, responsively to the record, a portion
5 of the data from the secondary storage subsystem to the primary storage subsystem for storage on the primary storage subsystem.

9. The method according to claim 1, wherein maintaining
10 and updating the record comprise marking respective bits in a bitmap corresponding to the locations to which the data are to be written on the first and second non-volatile storage media.

15 10. The method according to claim 1, wherein maintaining the record comprises maintaining a copy of the record on the primary storage subsystem, and wherein sending the message comprises deciding at the primary storage subsystem to send the message responsively to the copy of the record.

20

11. The method according to claim 10, wherein sending the message comprises modifying both the record and the copy of the record responsively to the specified location.

25 12. The method according to claim 11, wherein modifying both the record and the copy of the record comprises adding

a plurality of locations, including the specified location,
to both the record and the copy of the record.

13. The method according to claim 10, wherein maintaining
5 the copy of the record comprises selecting one or more
locations, other than the specified location, to be removed
from the record, and instructing the secondary storage
subsystem to remove the one or more locations from the
record, so as to limit a size of the record.

10

14. The method according to claim 13, wherein storing the
data comprises copying the data to be stored in the one or
more locations from the primary storage subsystem to the
secondary storage subsystem, and wherein selecting the one
15 or more locations comprises receiving a return message from
the secondary storage subsystem indicating that the
secondary storage subsystem has received the copied data,
and selecting the one or more locations to be removed from
the record responsively to the return message.

20

15. The method according to claim 13, wherein selecting the
one or more locations comprises identifying the locations at
which the first and second non-volatile storage media
contain substantially identical data, and selecting for
25 removal one of the identified locations that was least-
recently added to the record.

16. The method according to claim 13, wherein sending the message comprises adding one or more entries to both the record and the copy of the record responsively to the specified location, and grouping the entries added to the copy of the record and the record in generations according to an order of adding the entries to the records, and wherein selecting the one or more locations comprises determining at the primary subsystem that all the entries in one of the generations may be removed from the record.

10

17. The method according to claim 13, wherein instructing the secondary storage subsystem to remove the one or more locations comprises appending an instruction to the message sent from the primary storage subsystem to the secondary storage subsystem.

15

18. The method according to claim 1, wherein sending the message causes the secondary storage subsystem to predict one or more further locations to which the host processor is expected to write the data in a subsequent write operation, and to add the one or more further locations to the record.

20

19. The method according to claim 18, wherein the one or more further locations comprise a predetermined number of consecutive locations in proximity to the specified location.

25

20. The method according to claim 18, wherein maintaining the record comprises recording the locations to which the data are written using an object-based storage technique, and wherein the one or more further locations are chosen
5 based on a logical connection between storage objects.

21. A data storage system, comprising:

a primary storage subsystem, which comprises first non-volatile storage media; and

10 a secondary storage subsystem, which comprises second non-volatile storage media, and which is arranged to maintain a record that is predictive of locations to which data are to be written on the primary storage subsystem by a host processor,

15 wherein the primary storage subsystem is arranged to receive the data from a host processor for writing to a specified location, and to store the data in the specified location on the first non-volatile storage media while copying the data to the second storage subsystem, which is
20 arranged to store the data in the specified location on the second non-volatile storage media, and

wherein the primary storage subsystem is further arranged, upon receiving from the host processor the data to be written to a specified location on the first non-volatile
25 storage media, if the specified location is not included in the record, to send a message to the secondary storage subsystem so as to cause the secondary storage subsystem to update the record and to return an acknowledgment to the

primary storage subsystem indicating that the record has been updated, and

5 wherein the primary storage subsystem is further arranged to signal the host processor that the data have been stored in the data storage system responsively to receiving the data and, if the specified location was not included in the record, responsively to receiving the acknowledgment from the secondary storage subsystem.

10 22. The system according to claim 21, wherein the message sent to the secondary storage subsystem comprises the data, which are copied synchronously from the primary storage subsystem to the secondary storage subsystem.

15 23. The system according to claim 22, wherein the primary storage subsystem is arranged, if the specified location is included in the record, to copy the data from the primary storage subsystem to the secondary storage subsystem asynchronously, without causing the secondary storage
20 subsystem to update the record with respect to the specified location.

24. The system according to claim 23, wherein the first and second non-volatile storage media are located at mutually-
25 remote sites, and wherein at least one of the primary and secondary storage subsystems is arranged to transmit the data over a communication link between the sites.

25. The system according to claim 23, wherein the primary storage subsystem is arranged, if the specified location is included in the record, to signal to the host processor that the data have been stored without waiting to receive the
5 acknowledgment from the secondary storage subsystem.

26. The system according to claim 21, wherein the secondary storage subsystem is arranged to mirror the data held by the primary storage subsystem.

10

27. The system according to claim 26, wherein upon occurrence of a failure in the primary storage subsystem, the secondary storage subsystem is configurable to serve as the primary storage subsystem so as to receive further data
15 from the host processor to be stored by the data storage system.

28. The system according to claim 26, wherein upon recovery of the system from a failure of the primary storage
20 subsystem, the secondary storage subsystem is arranged to convey, responsively to the record, a portion of the data from the second non-volatile storage media to the primary storage subsystem for storage on the first non-volatile storage media.

25

29. The system according to claim 21, wherein the record comprises a bitmap, and wherein the secondary storage

subsystem is arranged to mark respective bits in the bitmap corresponding to the locations to which the data are to be written by the host processor.

5 30. The system according to claim 21, wherein the primary storage subsystem is arranged to maintain a copy of the record, and to determine whether to send the message responsively to the copy of the record.

10 31. The system according to claim 30, wherein the primary and secondary storage subsystems are arranged to update the copy of the record and the record, respectively, responsively to the specified location.

15 32. The system according to claim 31, wherein the primary and secondary storage subsystems are arranged to update the copy of the record and the record by adding a plurality of locations, including the specified location, to both the first and second records.

20

33. The system according to claim 30, wherein the primary storage subsystem is arranged to select one or more locations, other than the specified location, to be removed from the copy of the record, and to instruct the secondary
25 storage subsystem to remove the one or more locations from the record, so as to limit a size of the record.

34. The system according to claim 33, wherein the secondary storage subsystem is arranged to send a return message to the primary storage subsystem, indicating that the secondary storage subsystem has received the copied data, and wherein
5 the primary storage subsystem is arranged to select the one or more locations to be removed from the record responsively to receiving the return message.

35. The system according to claim 33, wherein the primary
10 storage subsystem is arranged to identify the locations at which the first and second non-volatile storage media contain substantially identical data, and to select for removal one of the identified locations that was least-recently added to the record.

15

36. The system according to claim 33, wherein the primary and secondary storage subsystems are arranged to respectively add one or more entries to both the copy of the record and the record responsively to the specified
20 location, and to group the entries added to the first and second records in generations according to an order of adding the entries to the records, and wherein the primary storage subsystem is arranged to determine that all the entries in one of the generations may be removed from the
25 record, and to instruct the secondary storage subsystem to remove all the entries in the one of the generations from the record.

37. The system according to claim 33, wherein the primary storage subsystem is arranged to append an instruction to the message sent to the secondary storage subsystem, so as to instruct the secondary storage subsystem to remove the
5 one or more locations from the record.

38. The system according to claim 21, wherein the message causes the secondary storage subsystem to predict one or more further locations to which the host processor is
10 expected to write the data in a subsequent write operation, and to add the one or more further locations to the record.

39. The system according to claim 38, wherein the one or more further locations comprise a predetermined number of
15 consecutive locations in proximity to the specified location.

40. The system according to claim 38, wherein the secondary storage subsystem is arranged to maintain the record using
20 an object-based storage technique, and to predict the one or more further locations based on a logical connection between storage objects.

41. A computer software product for use in a data storage
25 system including primary and secondary storage subsystems, which include respective first and second control units and respective first and second non-volatile storage media, the

product comprising a computer-readable medium in which
program instructions are stored, which instructions, when
read by the first and second control units, cause the first
control unit to receive data from a host processor for
5 writing to a specified location, and to store the data in
the specified location on the first non-volatile storage
media while copying the data to the second storage
subsystem, and cause the second control unit to maintain a
record that is predictive of locations to which the data are
10 to be written on the primary storage subsystem by the host
processor, and to store the data copied to the second
storage subsystem in the specified location on the second
non-volatile storage media,

wherein the instructions further cause the first
15 control unit, if the specified location is not included in
the record, to send a message to the secondary storage
subsystem so as to cause the second control unit to update
the record and to return an acknowledgment to the primary
storage subsystem, and cause the first control unit to
20 signal the host processor that the data have been stored in
the data storage product responsively to receiving the data
and, if the specified location was not included in the
record, responsively to receiving the acknowledgment from
the second control unit.

25

42. The product according to claim 41, wherein the message
sent to the secondary storage subsystem comprises the data,
which are copied synchronously from the primary storage
subsystem to the secondary storage subsystem.

43. The product according to claim 42, wherein the instructions cause the first control unit, if the specified location is included in the record, to copy the data from the primary storage subsystem to the secondary storage subsystem asynchronously, without causing the second control unit to update the record with respect to the specified location.

10 44. The product according to claim 43, wherein the first and second non-volatile storage media are located at mutually-remote sites, and wherein the instructions cause at least one of the first and second control units to transmit the data over a communication link between the sites.

15

45. The product according to claim 43, wherein the instructions cause the first control unit, if the specified location is included in the record, to signal to the host processor that the data have been stored without waiting to receive the acknowledgment from the second control unit.

20

46. The product according to claim 41, wherein the instructions cause the first and second control units to mirror the data held by the primary storage subsystem on the secondary storage subsystem.

25

47. The product according to claim 46, wherein the instructions cause the secondary storage subsystem, upon occurrence of a failure in the primary storage subsystem, to serve as the primary storage subsystem so as to receive
5 further data from the host processor to be stored by the data storage system.

48. The product according to claim 46, wherein upon recovery of the system from a failure of the primary storage
10 subsystem, the instructions cause the second control unit to convey, responsively to the record, a portion of the data from the second non-volatile storage media to the primary storage subsystem for storage on the first non-volatile storage media.

15

49. The product according to claim 41, wherein the record comprises a bitmap, and wherein the instructions cause the second control unit to mark respective bits in the bitmap corresponding to the locations to which the data are to be
20 written by the host processor.

50. The product according to claim 41, wherein the instructions cause the first control unit to maintain a copy of the record, and to determine whether to send the message
25 responsively to the copy of the record.

51. The product according to claim 50, wherein the instructions cause the first and second control units to update the copy of the record and the record, respectively, responsively to the specified location.

5

52. The product according to claim 51, wherein the instructions cause the first and second control units to update the copy of the record and the record by adding a plurality of locations, including the specified location, to
10 both the first and second records.

53. The product according to claim 50, wherein the instructions cause the first control unit to select one or more locations, other than the specified location, to be
15 removed from the copy of the record, and to instruct the second control unit to remove the one or more locations from the record, so as to limit a size of the record.

54. The product according to claim 53, wherein the
20 instructions cause the second control unit to send a return message to the primary storage subsystem, indicating that the secondary storage subsystem has received the copied data, and wherein the instructions cause the first control unit to select the one or more locations to be removed from
25 the record responsively to receiving the return message.

55. The product according to claim 53, wherein the instructions cause the first control unit to identify the locations at which the first and second non-volatile storage media contain substantially identical data, and to select
5 for removal one of the identified locations that was least-recently added to the record.

56. The product according to claim 53, wherein the instructions cause the first and second control units to
10 respectively add one or more entries to both the copy of the record and the record responsively to the specified location, and to group the entries added to the first and second records in generations according to an order of adding the entries to the records, and wherein the
15 instructions cause the first control unit to determine that all the entries in one of the generations may be removed from the records, and to instruct the second control unit to remove all the entries in the one of the generations from the record.

20

57. The product according to claim 53, wherein the instructions cause the first control unit to append an instruction to the message sent to the secondary storage subsystem, so as to instruct the second control unit to
25 remove the one or more locations from the record.

58. The product according to claim 41, wherein the instructions cause the second control unit, responsively to

the message, to predict one or more further locations to which the host processor is expected to write the data in a subsequent write operation, and to add the one or more further locations to the record.

5

59. The product according to claim 58, wherein the one or more further locations comprise a predetermined number of consecutive locations in proximity to the specified location.

10

60. The product according to claim 58, wherein the instructions cause the second control unit to maintain the record using an object-based storage technique, and to predict the one or more further locations based on a logical

15 connection between storage objects.